Marketing Science Institute Working Paper Series 2026

Report No. 26-111

Selling the Haggle

Keyan Li, Zelin Li, Siqi Pei and Feng Yang

Selling the Haggle

Keyan Li

Mendoza College of Business, University of Notre Dame

keyan.kli@nd.edu

Zelin Li

Sloan School of Management, Massachusetts Institute of Technology

jameszli@mit.edu

Siqi Pei

Eli Broad College of Business, Michigan State University

peisiqi@msu.edu

Feng Yang

College of Business, Shanghai University of Finance and Economics

yangf-df@dfl.com.cn

Date of submission:

March 9, 2026

Keyan Li is the corresponding author. He is an Assistant Professor of Marketing at the Mendoza College of Business, University of Notre Dame. Zelin Li is a Ph.D. Candidate in Marketing at the Sloan School of Management, Massachusetts Institute of Technology. Siqi Pei is an Assistant Professor of Marketing at the Eli Broad College of Business, Michigan State University. Feng Yang is a Ph.D. Candidate in Marketing at the College of Business, Shanghai University of Finance and Economics. The first three authors are listed alphabetically and contributed equally to this work. All remaining errors are the authors' own.

1 Introduction

Price negotiation is ubiquitous in personal selling – in automobiles, real estate, and B2B procurement, the final transaction price typically emerges from back-and-forth exchange. Yet industry and academia have argued that haggling is also a source of friction: customers face bargaining costs – hassle, discomfort, time – that deter many from negotiating, while firms bear substantial operational costs; eliminating negotiation benefits both sides (Scott Morton et al., 2011). Firms like Tesla have shifted toward fixed pricing. Besides customers' disutilities, some studies argue that the *outcomes* of haggling can be efficient. When customers' valuations are private information, firms may prefer negotiation as a segmentation or screening device to extract surplus and soften competition (Desai & Purohit, 2004). Both perspectives therefore evaluate negotiation through its price outcomes or the frictions it imposes.

We examine negotiation from a sales management perspective: does the *process* of negotiation affect whether a sale occurs? We study negotiation as a *sales device* – a tool the salesperson uses to close deals – and find that eliminating it can backfire. Using data from a major automaker and its dealership network in China that introduced a fixed-price policy for one of its car models, we find that removing negotiation *reduced* total sales volume. Although the fixed price was below the prior average negotiated price at most dealerships, the pattern persists within this subgroup. The effect cannot be explained by the price level alone; the missing piece is the salesforce and the selling process.

The conventional framework treats the sales encounter as a bilateral exchange between price-setting firms and a price-taking consumer. In markets organized around personal selling, however, the interaction is mediated by someone whose incentives may diverge from both: the salesforce. Existing theoretical studies that incorporate the salesforce treat negotiation as a device for the *principal*: firms use salesperson lobbying to design agents' incentives and elicit demand information (Simester & Zhang, 2014). Our argument is different. Negotiation also serves the *agent's own* interests as their deal-closing tool. Empirically isolating this role requires a setting in which negotiation is removed for some transactions but retained for others, and the negotiation process is at least partially observable.

We study such a setting using data from the leading automaker in China. In October 2024, the automaker introduced a one-price policy for the V0,[1] a mainstream model with an average transaction price of approximately 100,000 yuan (1 USD $\approx$ 7.1 yuan as of October 2024). The policy eliminated all vehicle price negotiation for V0 while retaining full negotiation for every other model sold at the same dealerships by the same salespeople. This within-dealer variation creates a quasi-experimental design – the same salesperson, serving the same local market, under the same management and compensation structure, simultaneously sells models with and without negotiation. We observe the universe of customer leads, call transcripts, and completed transactions across the automaker's entire dealership network, allowing us to track not only sales outcomes but also the customer journey – from initial inquiry through salesperson assignment and final conversion or attrition.

Our identification strategy uses synthetic difference-in-differences (SDID; Arkhangelsky et al., 2021), estimating the policy's causal effect at the dealer level. Because each salesperson sells multiple models at the same dealership, sales may spill over across models – making within-year cross-model comparisons unreliable. We therefore calendar-match each dealer to its 2023 observations, when no one-price policy existed, as the donor pool from which SDID constructs a synthetic counterfactual for the 2024 policy period. Year-over-year baselines are well established when within-period controls are contaminated.[2] SDID re-weights these donor units and pre-treatment time periods to match each treated dealer's pre-policy trajectory, constructing a counterfactual that accommodates level differences between the policy and pre-policy periods. We also use SDID with a donor pool of price-distant premium models as an alternative specification.

We find that the one-price policy *reduced* sales of the targeted model (V0) by 11.8 percent, driving a 3.25 percentage point drop in its share of total dealer sales; total dealer sales fell by 8.5 percent. In addition, the sales share of other car models increased by 2.23 percentage points. The decline persists even among dealers where the fixed price was lower than the pre-policy average transaction price (79% of the sample) – dealers where V0 effectively became cheaper.[3] The loss suggests a change in the negotiation *process* itself, as mediated by the salesperson.

[1] All model identifiers are anonymized.

[2] See, for example, Goldberg et al. (2024); Chetty et al. (2024); Chiou & Tucker (2022).

[3] The price-split test rules out the price level as the only driver but does not rule out all price channels. The loss of price *dispersion* – the ability to offer steep discounts to marginal customers – remains a plausible co-channel that our design cannot fully separate from the salesforce's behavioral response.

We argue that the mechanism is the removal of the salesforce's effective closing instrument. In personal selling, while the firm designs the pricing regime, the salesforce executes it, relying on negotiation as an effective instrument to convert a hesitant lead into a closed deal. A price alone does not close a sale – the salesperson must actively engage the customer. Price negotiation helps in this conversion process: it gives the salesperson a lever to address price resistance and keep hesitant customers in the conversation, increasing the likelihood of conversion. Importantly, the value of negotiation to salespeople operates through the sales *process* itself rather than only through the final price *outcome*.

Removing this instrument makes every salesperson less effective at closing deals on the fixed-price model, lowering the return to selling effort on it relative to negotiable alternatives. Because salespeople are compensated on conversion, they adapt by reallocating effort and persuasion – with the same customer – toward products where negotiation remains available. The most visible consequence is demand switching: customers intending to buy the fixed-price model are more likely to switch to the dealer's other models after the policy, consistent with salesforce-mediated redirection rather than a model-specific demand shock. Notably, Tesla – the most visible fixed-price automaker – complements its no-haggle policy with a different sales model: salaried advisors facilitate purchases rather than close deals on commission. Our setting isolates the removal of negotiation while the traditional commission-based salesforce remains in place.

The evidence is consistent with this salesforce-mediated mechanism: demand redirection is strongest where reliance on negotiation was highest before the policy. Dealers whose salespeople frequently initiated price discussions before the policy show a 6.6 percent increase in customer model-switching, compared with a 3.1 percent increase among low-reliance dealers. The effect concentrates among those who relied most on negotiation as a core element of their sales process. In addition, post-policy, the rate at which salespeople mention alternative models rises by 5.4 percentage points more at dealers with high reliance on price haggling than at low-reliance dealers. Lead quality metrics – conversion rates, decision duration, and store visit rates – show no significant deterioration, ruling out lead quality as an alternative explanation.

Our contributions are threefold. First, we show that the *process* of negotiation matters for whether sales occur. The existing literature studies negotiation through the firm's *external* factors: how bargaining affects customer utility, and how the resulting price dispersion enables firms to segment customers and extract surplus. We provide empirical evidence for an *internal* channel: the salesforce that executes the sales process. Using a natural experiment that removed negotiation for one model within a multi-model automaker, we find that sales decline even when the fixed price is set below the prior average, consistent with a salesforce-driven effect: removing negotiation degrades the agent's ability to convert leads into sales.

Second, we leverage within-dealer demand switching to identify the channel through which the salesforce and sales process matter. When negotiation is removed for one model but retained for others sold by the same salesperson, purchase composition shifts toward negotiable alternatives – and this reallocation concentrates among salespeople who relied more on price discussion before the policy. This is an agent-side incentive response: salespeople compensated on conversion redirect customers toward negotiable alternatives, where the closing tool remains effective. Our contribution is not just to estimate the switching effect *per se*, but to use this switching heterogeneity to distinguish salesforce-driven redirection from demand-side preferences.

Third, we argue that the friction we identify – degraded selling effectiveness when negotiation is removed – is not specific to the multi-product setting. Under a uniform fixed-price regime, the salesforce faces the same loss of its closing instrument with no margin for reallocation, making the friction worse but empirically unobservable. Our setting's asymmetric policy provides the variation necessary to isolate a new cost of fixed pricing: the salesforce friction.

The remainder of the paper is structured as follows. Section 2 reviews related literature and describes the institutional background of China's automobile retail market and the one-price policy. Section 3 describes the data and presents model-free evidence. Section 4 details the empirical strategy. Section 5 presents the main results and heterogeneous effects. Section 6 investigates mechanisms and presents robustness checks. Section 7 discusses implications and concludes.

2 Literature Review and Institutional Background

2.1 Related Literature

Our paper contributes to three branches of literature: posted-price versus bargaining mechanisms, salesforce management, and empirical automobile pricing.

Posted Price vs. Bargaining. A foundational question in market design is when a seller should post a price versus negotiate. Riley & Zeckhauser (1983) show that a posted price is the optimal selling mechanism. Empirically, bargaining imposes costs on customers (Scott Morton et al., 2011); it produces discriminatory outcomes (Ayres & Siegelman, 1995); and a substantial share of bargaining pairs fail to trade (Larsen, 2021). Some of this work engages with the haggling process, but frames it as a friction for customers. A parallel strand argues that negotiation *outcomes* can be efficient for firms, when under private information (Myerson & Satterthwaite, 1983) or when customer valuations are heterogeneous (Wang, 1995): haggling serves as a valuation screening device and extracts more surplus (Desai & Purohit, 2004), and resolves moral hazard by tying the price to observed quality (Bester, 1993; Gill & Thanassoulis, 2016). The case for price negotiation rests primarily on its pricing outcomes, *external* to firms. We pose a distinct empirical question: does negotiation generate value to firms through the *selling process* itself – through the salesperson's use of haggling as a closing tool? This *internal* salesforce dimension is overlooked in existing empirical studies. We shift the lens from mechanism design to sales management and incentive design.

Salesforce Management. While no prior work connects negotiation to the salesperson's selling process, a broader literature studies how salespeople respond to organizational design. Early conceptual frameworks characterize selling as a contingent, adaptive process (Weitz, 1981; Weitz et al., 1986; Anderson & Oliver, 1987). More recent empirical work shows that salespeople re-optimize effort when firm design changes: Chung et al. (2014) show that bonus structures alter the timing and intensity of effort; Yang et al. (2019) estimate that salespeople steer customers toward higher-margin products; Li et al. (2020) show that team structures shift effort allocation across tasks, while Lim & Chen (2014) show that group incentives change how salespeople exert effort.

Salesforce Management and Pricing. Existing work (primarily theoretical) studying negotiation through the internal salesforce lens focuses on pricing delegation and principal-agent screening. Lal (1986) shows that delegation helps when the agent possesses private market information. Joseph (2001) identifies that limiting the salesperson's pricing discretion can force effort toward high-valuation customers rather than substituting effort with price concessions. Simester & Zhang (2014) show that firms use salespeople's lobbying for price concessions to extract private demand information, and negotiation serves the principal's screening needs. These models treat negotiation as a mechanism the firm designs. We study empirically what happens when negotiation is removed as the salesperson's closing tool. The salesperson's incentivized response represents a different principal-agent friction.

Pricing in Automobile Markets. The automobile market is a canonical empirical setting for studying price negotiation. Prior work documents discriminatory outcomes in negotiated transactions across customer groups (Ayres & Siegelman, 1995; Goldberg, 1996; Scott Morton et al., 2003). Other studies examine how asymmetric information (Scott Morton et al., 2001; Zettelmeyer et al., 2006), brand preferences (Chen et al., 2008), and competition (Zeng et al., 2016) shape negotiated price outcomes. Jindal & Newberry (2018) estimate customers' bargaining costs and argue that these costs enable second-degree price discrimination by segmenting the market. Two papers study what happens when negotiation is removed. Zeng et al. (2016) find that Toyota Canada's brand-wide no-haggle policy raised prices by segmenting no-hagglers and softening cross-brand competition. Busse et al. (2010) show that when automakers replaced negotiated pricing with employee pricing, many customers paid more because posted prices shifted beliefs about current prices relative to future deals. Across these studies, the focus is on customer beliefs or equilibrium pricing outcomes, with limited attention to the sales process that generates prices. Our paper provides direct causal evidence that eliminating haggling reduces sales through a salesforce channel, specifically the salesperson's behavioral response to losing a key closing instrument.

2.2 Institutional Background

2.2.1 China's Automobile Retail Market

China is the world's largest automobile market by volume, with annual new passenger vehicle sales exceeding 26 million units.[4] Retail distribution is organized primarily around the *4S dealership model* – franchised, brand-exclusive outlets that integrate Sales, Service, Spare parts, and Surveys under a single roof. Each dealer holds a franchise agreement with a single manufacturer or joint-venture (JV) partner and serves as the brand's local retailer. The manufacturer sets a manufacturer's suggested retail price (MSRP) for each model, but this MSRP serves as a starting point for negotiation rather than the transaction price. The majority of passenger vehicles continue to be sold through franchised 4S dealers where prices are negotiated.[5]

In this setting, the typical car purchase involves substantial price negotiation. A prospective buyer visits dealerships, discusses pricing and terms with a salesperson, and arrives at a final transaction price through back-and-forth exchange. The result is considerable price opacity – the same model at the same brand can transact at different prices across dealers and cities. Critically for our study, the salesperson is the central intermediary in this process, exercising meaningful discretion over the size and structure of the deal.

Several JV brands have experimented with one-price policies in recent years. The firm's one-price policy, which we study, emerged within this broader trend.

2.2.2 The One-Price Policy

The focal automaker operates one of China's most extensive dealer networks, with over 1,100 retail points nationwide selling passenger vehicles. On October 20, 2024, the firm launched the first model under its one-price initiative (V0) via a nationwide livestream event, introducing a one-price – a fixed, non-negotiable transaction price for the vehicle – at 99,800 yuan (MSRP: 125,900 yuan). Critically for our identification strategy, all other models retained full price negotiation at the same dealerships, sold by the same salespeople under the same compensation

[4] China Association of Automobile Manufacturers (CAAM) annual statistics. Passenger vehicle sales reached 26.06 million in 2023 and approximately 28.6 million in 2024.

[5] As of 2024, direct-sales and agency models account for a growing but still minority share of total passenger vehicle retail.

structure.

The firm's stated rationale [6] for the policy centered on consumer experience: eliminating the anxiety associated with dealership negotiations ("no tricks, no time pressure"), providing price transparency and consumer trust, and shifting competitive emphasis from price concessions to service quality. The policy also aimed to unify pricing across the dealer network and to appeal to younger consumers accustomed to the posted-price experience of direct-sales brands.This rationale aligns with the conventional wisdom that negotiation is pure friction: if customers dislike haggling, removing it should increase demand.

In operational terms, the one-price is the vehicle's *final transaction price* – the amount invoiced to the customer with zero room for salesperson concession. The policy is enforced through the firm's central invoicing system, which requires dealers to issue invoices at the posted one-price. This eliminates sticker-price negotiation, add-on bundling discounts, and "friends and family" adjustments. To contextualize the price levels, 79% of dealers had a pre-policy average V0 transaction price above the one-price threshold, while only 21% were below it. The one-price thus represented an effective discount for the majority of dealers' customer base. The firm did not alter the salesperson compensation structure during our study window. Because the firm subsequently extended the one-price to S2 on November 22, 2024 (Figure 1),[7] our main analysis focuses on the initial V0 policy window (October 20–November 20, 2024), when V0 was the only model under fixed pricing. We exclude an anticipation window (the week prior to the policy's onset) in both years to guard against anticipation effects. As a robustness check, Table B.4 in the appendix reports that shifting the treatment date one week earlier to include this window as treated (Roth et al., 2023) yields slightly stronger estimates. [8]

[6] This rationale is based on our own extensive interviews with the firm's executives.

[7] Tables and Figures follow References throughout.

[8] Because the policy was announced on October 20 with dealer-facing communications shortly beforehand, the calendar week spanning October 13–19 is a transition period that may contain partial exposure or anticipatory behavior. Canonical DiD identification relies on a no-anticipation assumption (Roth et al., 2023); rather than force this week into either the pre or post period, we exclude it from both years symmetrically – analogous to the "donut" designs used in regression discontinuity when the cutoff neighborhood is contaminated (Barreca et al., 2011). Including the anticipation week would, if anything, attenuate the estimate by mixing treated and untreated days within the same observation, so the exclusion is conservative. Table B.4 shifts the treatment date one week earlier to include the anticipation window as treated.

3 Data

3.1 Data Sources

We use proprietary administrative data from the internal customer relationship management (CRM) and sales management systems of one of the top automotive manufacturers in China. The data cover the universe of customer interactions and completed transactions across the firm's nationwide dealer network. Multiple linked datasets track the full customer journey – from initial lead generation and dispatch through follow-up contact, order placement, vehicle delivery, and invoicing – spanning January 2023 to May 2025, with coverage varying by dataset. We describe each dataset in turn.

Lead Data. The lead management dataset contains approximately 34.1 million lead-action observations. The variables cover dealer information, customer attributes, lead characteristics, the customer's intended vehicle, the lead lifecycle, salesperson assignment, and conversion outcomes. Lead sources are predominantly digital: approximately 88.7% originate from online channels, with dealer-generated and dealer social media channels accounting for the balance.

Customer Follow-up Records. The follow-up detail dataset comprises approximately 2.04 million unique customer-by-record-creation-time observations, constructed by combining standard follow-up logs with call records and deduplicating on customer and follow-up time. Critically, a subset of these records includes verbatim call transcripts (N = 882,952). Beginning in 2024, the company adopted virtual phone numbers, and calls placed using these virtual numbers have associated transcripts. The transcripts enable textual analysis of whether and how salespeople discuss pricing, a key input for our mechanism analysis in Section 6.

Transaction Data. The car delivery dataset records approximately 1,336,114 completed delivery records (i.e., individual transactions) and 1,276,069 invoice records, with information including car models, customer demographics, identifiers, order time, assigned salesperson, the vehicle's base price, finalized transaction price, payment method, and order channel. The dataset uniquely links to invoice records via the vehicle identification number (VIN).

3.2 Sample Construction and Summary Statistics

Our analysis sample is constructed by linking the transaction and lead datasets through a three-step record linkage pipeline. First, we merge invoice records with car delivery records on VIN. After resolving duplicate VINs in car delivery records, we achieve a 94% match rate (1,195,531 of 1,276,069 invoice records). We drop 19 observations with negative transaction prices. Unmatched records are predominantly company purchases. Second, we merge the resulting invoice-delivery dataset with lead data, yielding 466,612 unique matched orders – representing 39.3% of invoice-delivery records and 1.37% of total leads, consistent with the fact that most leads do not convert to purchases. Unmatched invoice-delivery records are largely walk-in or referral sales not captured in the lead system and exhibit a higher rate of company purchases (2.92% vs. 1.16% among matched records). Third, we merge the combined dataset with follow-up records on the customer record creation time and customer name. This yields 295,951 leads with matched follow-up information (78.77% match rate), restricting the sample to leads whose lead creation date or invoice date falls in the analysis window.

We impose several sample restrictions to ensure a coherent sample. We require that lead creation time precedes the invoice date, dropping observations with missing or inconsistent timestamps.[9] We exclude returns and non-individual purchases identified by VAT special invoices. We then drop three premium car series whose average invoice prices lie far above V0's price range: P1 (approximately 230,000 yuan), P2 (approximately 176,000 yuan), and P3 (approximately 172,000 yuan). These are substantially more expensive than V0 and exhibit minimal switching among leads initially intending to purchase V0 (1.2%, 1.1%, and 0.4%, respectively). We exploit these premium models as an alternative control group for estimating the effect of the one-price policy in a robustness specification. The remaining eight series, listed in Table 1, cluster in a 70,000–137,000 yuan band around V0's approximately 101,000 yuan.

We construct a panel by stacking each dealer's 2024 observations (treated) with its calendar-matched 2023 observations (donor). The policy date anchor is October 20 in each year, with weekly frequency. As described in Section 2.2.2 on our analysis window construction, we drop a seven-day anticipation window (October 13–19) in both years, shifting pre-treatment

[9] The median invoice-order date gap is 0 days (mean = 4.0 days).

periods by one week to close the gap (Figure 1). The resulting panel comprises 688 treated dealer-units (2024) and 795 donor dealer-units (2023). Table B.4 in the appendix confirms that including the anticipation week as treated yields slightly stronger estimates.

Table 2 presents summary statistics for the estimation sample. Panel A reports transaction outcomes: total sales counts, V0 and other model sales separately, and V0's share of total dealer sales. Panel B reports lead and funnel metrics. Sales volumes, lead counts, and decision durations are rescaled by constant factors for confidentiality; shares and rates are reported in natural units.

As model-free evidence, we examine customer model-switching patterns by comparing intended and actual purchased models before and after the policy (Figure 2). Before the policy, 29.8% of V0-intended customers ended up purchasing a different model; after the policy, this rate rose to 40.6% – a significant increase of 10.8 percentage points. By contrast, the change in switching rate among customers intending to buy other models was -0.6 percentage points and statistically insignificant.

4 Empirical Strategy

4.1 Conceptual Framework

As we discuss in the Introduction, negotiation in personal selling is not only a way to set prices; it is also part of the interaction sales process before closing the deal. The salesperson must invest effort to convert a customer's interest into commitment – discovering needs, framing the product, handling objections, and negotiating terms. Price negotiation matters to this process: it gives the salesperson a concrete, adaptable response to price resistance, builds commitment through back-and-forth exchange, and keeps hesitant customers engaged. Salespeople who rely heavily on price discussion as a closing device – and customers who depend on guided selling – are especially exposed when this instrument is not available.

Our setting eliminates negotiation flexibility for one product while retaining it for every other product sold by the same salesperson at the same firm. This creates a natural margin for cross-product substitution. If negotiation operates as a closing tool, weakening it hurts the salesperson's ability to close the deal with the current lead. Because salespeople are compensated on conversion, they are incentivized to redirect the same lead toward a product where negotiation remains available – reallocating effort from a harder close to an easier one. This reallocation incentive is a form of multi-task moral hazard: when the return to effort on one task falls, agents shift effort toward tasks where returns remain higher (Holmström & Milgrom, 1991). In particular, the incentive is stronger for salespeople who relied most heavily on price discussion before the policy (and, symmetrically, for customers who are more dependent on guided selling), because the loss of negotiation is most binding in those interactions.

This logic yields five testable implications. First, V0 sales should fall after the policy (Section 5.1). Second, other models' sales share and V0-to-other switching should rise as salespeople redirect customers toward negotiable alternatives (Sections 5.1 and 6.1). Third, this demand switching should be strongest among leads handled by salespeople who relied most heavily on price negotiation before the policy, because the loss of selling effectiveness is largest for those who depended most on negotiation (Section 6.1). Fourth, the effect should be larger among customers who depend more on guided selling, such as older or less experienced buyers, because these customers are more susceptible to salesperson redirection (Section 5.2). Finally, within-dealer substitution need not preserve total sales: although fixed pricing may attract additional leads, the substitute model could be a worse match for the redirected customer's needs; as a result, some customers complete the purchase, whereas others delay or exit, and total sales volume may rise or fall (Section 5.1).

Some alternative accounts could explain a post-policy sales decline, and each generates a distinct empirical prediction. First, if the decline is driven by the *fixed-price level* being set too high and the standard price – sales relationship, the effect should disappear at dealers where the new fixed price is *cheaper* than the pre-policy average transaction price. If the mechanism is negotiation degraded as a selling tool, both groups may show similar declines, because the closing instrument is removed regardless of the price level. We test this directly in a price-level split (Section 6.2).

Second, if the decline reflects *uniform consumer discomfort* with fixed pricing – a demand-side preference for negotiation and getting "deals" – the effect should be homogeneous across salespeople with different pre-policy selling strategies. If the mechanism is tool removal, salespeople who relied most heavily on price discussion should show larger effects, because they lose the instrument they depended on most. We test this using pre-policy measures of salesperson price-discussion reliance (Section 6.1).

We note that our SDID design identifies the causal effect of the policy on sales, and the mechanism evidence – switching patterns and salesperson heterogeneity – is consistent with a sales-process channel but does not rule out all alternatives. We return to these distinctions when interpreting the results. We now turn to our identification strategy.

4.2 Identification Strategy

We estimate SDID (Arkhangelsky et al., 2021) at the dealer level, using each dealer's calendar-matched 2023 observations as the donor pool to construct the counterfactual trajectory. Each dealer observed in the 2024 policy window is treated; the same dealer's observations from the corresponding 2023 calendar window – when no one-price policy existed – serve as donors. SDID reweights the donor pool and pre-treatment time periods so that the synthetic counterfactual closely tracks each treated dealer's pre-policy trajectory. Because the 2023 donors are unaffected by the 2024 policy, the design avoids within-year spillover concerns that would arise if other 2024 models served as controls.

We organize the data into an $N \times T$ panel matrix $\mathbf{Y}$, where rows index dealer-year units and columns index time periods. Each dealer appears twice: as a 2023 donor unit and as a 2024 treated unit. Of the N units, N_{co} are donor units (2023 observations) and N_{tr} are treated units (2024 observations); of the T periods, T_{pre} are pre-treatment periods and T_{post} are post-treatment periods. The SDID estimator solves a weighted two-way fixed effects problem:

$$\left(\hat{\tau}^{\text{sdid}}, \hat{\mu}, \hat{\alpha}, \hat{\beta}\right) = \arg\min_{\tau,\mu,\alpha_i,\beta_t} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\,\tau\right)^2 \hat{\omega}_i\, \hat{\lambda}_t \right\}, \tag{4.1}$$

where Y_{it} is the outcome for unit i in period t, $W_{it} = \mathbb{1}[\text{unit } i \text{ is treated}] \times \mathbb{1}[t > T_{\text{pre}}]$ is the treatment indicator, α_i and β_t are unit and time fixed effects, and $\hat{\tau}^{\text{sdid}}$ is the estimated average treatment effect on the treated. Two sets of data-driven weights distinguish SDID from standard DiD and synthetic control. Unit weights $(\hat{\omega}_i)$ reweight 2023 donors so that their pre-treatment *trends* (not levels) match the treated group's trajectory; an intercept allows level differences, so the method matches on changes rather than on levels. Time weights $(\hat{\lambda}_t)$ reweight pre-treatment periods to emphasize those most informative for predicting post-treatment outcomes. We compute standard errors via the jackknife (Arkhangelsky et al., 2021). Table B.1 in the appendix reports the full weight vectors. The unit weights are highly dispersed across the 795 donor dealers, and the synthetic control draws broadly on the donor pool rather than relying on a few units.

4.2.1 Justification of Identification

We cannot use other 2024 models at the same dealers as controls. The demand switching we study directly contaminates them: customers intending to buy V0 may switch to alternative models, inflating those models' sales and violating the stable unit treatment value assumption. This spillover is not hypothetical – we document it empirically in the switching analysis (Section 6.1). Nor can we compare V0 across dealers: the policy applies nationwide, so no untreated V0 dealers exist in 2024. We therefore construct the counterfactual from the same dealers in the same calendar window one year earlier, when no one-price policy existed. Year-over-year baselines are well established in settings where within-period controls are contaminated by the treatment itself (Goldberg et al., 2024; Chetty et al., 2024; Chiou & Tucker, 2022).[10] A simple year-over-year difference, however, would be biased if 2024 experiences idiosyncratic shocks unrelated to the policy. SDID addresses this by reweighting the 2023 donor pool to match treated dealers' pre-policy *trajectories*, accommodating level differences between years while preserving the clean identification the YoY design provides.

We estimate an alternative SDID design that bypasses year-on-year comparisons and avoids relying on 2023 calendar-matched donors. Although cross-model comparisons within 2024 are generally contaminated, a subset of models may be largely insulated from this channel. Three premium models (P1, P2, and P3) are priced 1.7–2.3 times above V0, and only 2.7% of V0 customers switched to a premium model throughout the data period. We exploit this separation by estimating a within-year SDID at the dealer $\times$ model level, using these premiums as controls and avoiding cross-year comparisons entirely.

We conduct a series of validation tests. First, we assess the SDID counterfactual by checking

[10] Goldberg et al. (2024) use the prior year (pre-regulation) as a control group when evaluating the GDPR, which affected all EU websites simultaneously; Chetty et al. (2024) use same-week-2019 spending as the baseline for 2020 COVID impacts when the shock affected all sectors simultaneously; and Chiou & Tucker (2022) validate their design using year-over-year comparisons when evaluating an advertising restriction that affected behavior across categories.

pre-treatment fit. Figure 3 plots the treated dealers' outcome against the SDID synthetic control; the two series track closely in the pre-treatment period, with a root mean squared error (RMSE) of 0.102 for total sales and 0.056 for V0 sales (in log units). Second, we implement two placebo tests. The first is the standard in-space placebo for synthetic control designs (Abadie et al., 2010; Abadie, 2021): we randomly reassign treatment and donor labels and re-estimate SDID 500 times.[11] All outcomes fall in the extreme tail of the permutation distribution, significant at the 5% level (Table B.2), confirming that the effects are specific to the treated units rather than an artifact of cross-sectional structure. The second is an in-time placebo. A natural concern with the year-over-year design is that V0-specific seasonal factors in the October–November window, rather than the policy, drive the results. We run the identical specification in the same calendar window in 2023 using a placebo policy date of October 20, 2023, with the within-year premium-control SDID design: same window, same dealers, same control models, but no one-price policy. The placebo estimate is a statistically insignificant -0.004 (s.e. $= 0.020$, $p = 0.847$), suggesting the 2024 effect is not a seasonal artifact. Finally, we complement the dealer- and salesperson-level SDID with a transaction-level switching analysis (Section 6.1), which examines demand redirection using an interrupted time series logistic regression. This analysis is correlational: each customer appears only once, precluding a unit-level counterfactual over time.

5 Results

5.1 Effects on Sales

We begin by assessing the SDID counterfactual through its pre-treatment fit. Figure 3 plots the outcome trajectory for treated dealers alongside the corresponding SDID synthetic control. The two trajectories are closely aligned during the pre-treatment period, indicating that the optimally weighted counterfactual accurately tracks the treated units before the policy (pre-treatment RMSE = 0.056 for V0 sales in log units). After the policy date, the trajectories diverge sharply,

[11] The test evaluates whether the actual estimate is distinguishable from effects generated by random treatment assignments. If the estimated effect is genuine, few or no placebo reassignments should produce an estimate as large as the actual one. The permutation p-value measures the share that do; a value smaller than 0.05 indicates the effect is unlikely to arise by chance, and the design passes the placebo.

showing a divergence consistent with a discrete structural break at the one-price introduction.
The first column of Table 3 presents the SDID estimates from Equation 4.1. The one-price policy reduced V0 sales by 11.8 percent.[12] An 11.8% decline in V0 sales corresponds to approximately 0.15 fewer vehicles per dealer per week, evaluated at the rescaled sample mean of 1.29 V0 sales per dealer. Total dealer sales dropped by a statistically significant 8.5 percent. The policy also shifted the composition of dealer sales: V0's share of total dealer sales fell by 3.25 percentage points, while the share of other models rose by 2.23 percentage points. In terms of intention changes, we define V0-to-other switching as a binary indicator equal to one when a customer's intended model is V0 but the purchased model is not (see Appendix A for details). V0-to-other switching increases by 5.0%.

The second column of Table 3 re-estimates the SDID specification using the individual salesperson as the cross-sectional unit. The results remain qualitatively consistent. V0 sales decline by 3.0%. In terms of the composition of sales, V0's share decreases by 1.06 percentage points, while the sales share of other-model increases by 3.30 percentage points. In terms of intention changes, V0-to-other switching increases by 1.28%. All above estimates are significant at the 1% level. The effect on total sales is statistically insignificant at the salesperson level.

Our year-on-year SDID uses 2023 calendar-matched donors. As an alternative that avoids cross-year comparisons entirely, we estimate a within-year SDID using premium models (P1, P2, P3) as controls – models priced 1.7–2.3$\times$ above V0 that receive fewer than 3% of V0's pre-policy switching traffic. At the dealer $\times$ model level, V0 sales fall by 9.3%, significant at the 1% level, and V0-to-other switching increases by 8.5% – all consistent with the main year-over-year SDID estimates. Table B.3 and Figure B.2 in the appendix report the full results.

Before examining heterogeneous effects, we preview the alternative explanations ruled out in Section 6.2. Most importantly, the sales decline does not disappear at dealers where the fixed price was set below the prior average – both groups show symmetric V0 declines ($p = 0.734$), ruling out a price-level explanation. We also find no evidence that lead quality deterioration drove the decline. These robustness checks establish that the mechanism operates through the negotiation process itself, not price levels or demand quality.

[12] Throughout the paper, percentage changes are computed as $100 \times (\exp(\hat{\tau}) - 1)$, where $\hat{\tau}$ is the SDID estimate from the log-transformed outcome.

5.2 Heterogeneous Effects

We exploit cross-sectional heterogeneity in customer characteristics and salesperson behavior to further discriminate between the tool-removal mechanism and alternatives.

Table 4 (Panel A) splits customers into three age groups: 33 and below, 34–49, and 50 and above. The switching effect is monotonically increasing in age. Among the youngest group, the increase in V0-to-other switching is a statistically insignificant 0.7%. Among middle-aged customers (34–49), it rises to 2.9%, significant at the 5% level. Among customers aged 50 and above, the effect is largest at 4.1%, significant at the 1% level. The other-model sales share shows the same gradient: -0.38 percentage points for the youngest (insignificant), 3.95 percentage points for middle-aged, and 6.95 percentage points for the oldest, both significant at the 1% level. Older customers show the largest switching effects – the group that conventional wisdom predicted would benefit most from the simplicity of fixed pricing. The pattern is consistent with our mechanism: older customers tend to depend more on the salesperson's guidance during the purchase process, and when the salesperson's toolkit is constrained, they are the first to be redirected.

Panel B of Table 4 splits customers by prior purchase experience, classified via LLM analysis of conversation transcripts (see Appendix A for details). Among experienced buyers, V0-to-other switching increased by a statistically insignificant 0.5%, while among inexperienced buyers the increase was 1.6%, significant at the 10% level. The other-model sales share tells a consistent story: a statistically insignificant 2.55 percentage points for experienced customers versus 4.87 percentage points for inexperienced ones, significant at the 5% level. Inexperienced customers are more affected – again consistent with our mechanism that these customers depend more on salespeople to navigate the purchase process.

6 Mechanisms and Robustness

Section 5 established that the one-price policy reduced sales and shifted demand composition. We now investigate the mechanisms: we document behavioral patterns consistent with demand switching toward models where negotiation remained available, and that eliminating price flexibility degraded the broader sales process (Section 6.1). We close with robustness checks.

6.1 Mechanism: Lead Redirection

6.1.1 Demand Switching

We first examine whether demand shifted away from V0 in patterns consistent with salesperson-mediated switching. At the dealer level (Table 5), the count of V0-to-other switchers rose significantly by 5.0% after the policy (also reported in Table 3). Walk-in customers (those without a prior lead) who ended up purchasing non-V0 models also rose significantly by 3.1%. Measured as the other-model sales share conditional on customer intention, the pattern is consistent: among V0-intended customers, the other-model sales share increased significantly by 5.42 percentage points; among customers with no recorded intention, it rose significantly by 2.07 percentage points. In contrast, among customers with other model intention, the switching rate is an insignificantly increase. V0-intended customers shifted toward models where negotiation remained available, a pattern consistent with salesperson-mediated redirection.

As an alternative, we perform the demand switching analysis at the individual transaction level. The switching specification is:

$$\begin{aligned} \Pr(\text{Purchase}_i \neq \text{V0}) = \Lambda\big(&\beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Intended V0}_i + \beta_3 (\text{Post}_t \times \text{Intended V0}_i) \\ &+ \beta_4 \text{Walk-in}_i + \beta_5 (\text{Post}_t \times \text{Walk-in}_i) + \gamma_d\big), \end{aligned} \tag{6.1}$$

where i indexes the individual transaction, $\Lambda(\cdot)$ denotes the logistic function, Intended V0_i equals one if the customer's lead recorded V0 as the intended model, Walk-in_i indicates walk-in customers without a prior lead, and γ_d denotes dealer fixed effects. The coefficient β_3 captures whether V0-intended customers are more likely to purchase a non-V0 model after the policy. Because each customer only appears once – with one intention – in the dataset, all V0-intended customers are simultaneously affected. However, the SDID design requires the same units observed over time and can contribute to both treated and donor groups, which is possible for dealers and salespeople but not for individual customers. Therefore, this specification is an interrupted time series for V0 intenders rather than a SDID design, and we interpret these transaction-level estimates as suggestive correlational evidence complementing the SDID findings.

At the transaction level (Table 6), the demand switching analysis is consistent with the SDID pattern. After the policy, V0-intended customers were more likely to purchase other models, with an odds ratio of 1.811, significant at the 1% level. Walk-in customers without a prior lead also showed a significant increase in switching (odds ratio = 1.164, significant at the 1% level). To quantify the economic magnitude of the composition shift, we compute average marginal effects (AMEs) from the logistic regression.[13] Among customers who completed a purchase, the AME for V0-intended buyers is 12.7 percentage points, significant at the 1% level, indicating that the policy shifted the composition of purchases away from V0. The switching is concentrated among exactly those customers whose intended model lost negotiation.

6.1.2 Behind Demand Switching: Salesperson Reliance on Price Negotiation

We test the tool-removal mechanism through heterogeneity across salespeople (Table 7). After the policy, the rate at which salespeople brought up pricing as the first topic fell significantly by 2.2 percentage points, suggesting a general shift away from price-led conversations. This behavioral change should matter most where salespeople relied most heavily on price discussion. We classify dealers by their pre-policy average salesperson price-first discussion rate – the share of conversations in which the salesperson initiated the price topic, derived from textual analysis of call transcripts (see Appendix A for details) – using a median split (median = 48.6%; "high-reliance" vs. "low-reliance").

High-reliance dealers show a significant 6.6% increase in V0-to-other switching, roughly double the significant 3.1% increase among low-reliance dealers. The balance check (Table B.6 in the appendix) confirms the classification: high-reliance dealers' salespeople had a pre-policy price-first rate of 63.9%, compared to 32.2% for low-reliance dealers, while observable covariates, total conversations and number of salespeople per dealer, are balanced. This is consistent with the mechanism: the policy switches sales most where salespeople relied most heavily on price negotiation as a closing tool. Table B.7 in the appendix shows qualitatively similar results under two alternative definitions of reliance: (i) a dealer-level comparison restricted to the top and

[13]Conditional on completing a purchase, the AME for intention group g measures the average change in the predicted probability of purchasing a non-V0 model: $\widehat{\text{AME}}_g = \frac{1}{N_g}\sum_{i\in g}\left[\hat{p}_i(\text{post} = 1) - \hat{p}_i(\text{post} = 0)\right]$, where $\hat{p}_i(\text{post} = j) = \Lambda(\hat{\alpha}_{d(i)} + \hat{\beta}_1 j + \hat{\boldsymbol{\beta}}_2' \, \text{intended}_i + \hat{\boldsymbol{\beta}}_3' \, (j \times \text{intended}_i))$, $\Lambda(\cdot)$ is the logistic function, and $\hat{\alpha}_{d(i)}$ is the dealer fixed effect. Standard errors are computed via dealer-clustered bootstrap.

bottom quartiles of the price-first-rate distribution, and (ii) a salesperson-level classification based on salesperson-specific price-first rates.

A complementary measure, the ratio of salesperson price mentions to customer price mentions, constructed via keyword matching in call transcripts (see Appendix A for details), tells the same story (Table 7, Panel B). Dealers where salespeople exhibited a higher mention ratio show a 5.22% increase in V0-to-other switching, significant at the 5% level, while the lower-ratio group shows a significant but smaller 3.91% increase. Both measures of price-negotiation reliance yield similar pattern.

When Do Salespeople Mention Alternatives? We now look for direct evidence for salespeople mentioning alternative models in individual sales conversations. Using the same call transcripts, we construct two binary indicators at the lead level (see Appendix A for details): whether the salesperson uses *haggle-refusal language*, defined as phrases refusing to negotiate on V0's price, and whether the salesperson mentions another car model by name or through redirection phrases (e.g., "look at another one").

Table 8 reports results from 8,720 V0-intended leads with phone transcripts. Panel A tests whether the policy's effect on the other-model mention rate varies with salesperson reliance on price negotiation. At high-reliance dealers that relied more heavily on the haggling tool prior to the policy, the other-model mention rate increased by 3.1 percentage points after the policy, significant at the 5% level. In contrast, at low-reliance dealers, the change is statistically insignificant -1.5 percentage points.

Panel B reports estimates from a fixed effects specification; the results remain consistent with the previous pattern: the interaction term (Post $\times$ High reliance) is 5.4 percentage points and is significant at the 5% level. This result means that dealers whose salespeople relied more on price discussions pre-policy experienced a relative increase of 5.4 percentage points in the likelihood of a salesperson mentioning an alternative model. This indicates that when the price-haggling tools weakens, salespeople tend to redirect customers to different models.

6.2 Robustness Checks

We subject the main SDID findings to a set of robustness tests.

Ruling Out the Price-Level Explanation. Table 9 tests the most immediate alternative explanation – that the one-price was simply set too high, deterring price-sensitive customers. We split dealers according to whether their pre-policy average V0 final user price was above or below the one-price. Among dealers in our sample, 79% had average pre-policy final user prices above the one-price, meaning V0 became *cheaper* after the policy; the remaining 21% had average prices below, meaning V0 became *more expensive*. If the price level were driving the decline, the effect should concentrate among dealers for whom the one-price was an increase, and should disappear (or reverse) among dealers for whom it was a decrease. The results reject this prediction. Among dealers where V0 became cheaper (pre-policy average above the one-price), V0 sales fell significantly by 13.1% and total sales by 8.4%, significant at the 5% level. Among dealers where V0 became more expensive (pre-policy average below the one-price), V0 sales fell significantly by 14.8% but total sales declined by a statistically insignificant 10.1%. A formal test fails to reject the null that V0 sales effects are equal across price groups ($p = 0.734$), confirming symmetric decline. The sales decline operates through the removal of the negotiation process itself.

Lead Quality Falsification. A natural concern is that the one-price policy degraded the quality of incoming leads, independently explaining the sales decline. Table 10 tests this directly: conversion rates, decision duration, and store visit rates are all statistically insignificant for both V0-intended and non-V0-intended leads. These results provide no evidence that lead quality deterioration drove the sales decline, consistent with the policy's primary channel operating through salesperson behavior at the point of sale.

Window Sensitivity. Table B.5 in the appendix re-estimates the main SDID specification for V0 using different pre- and post-treatment windows around the policy date. The estimated effect is consistently negative across all window choices, ranging from -5.1% to -11.5%. The main specification qualitatively matching these estimates. Statistical significance increases with window length, as expected given that shorter windows provide fewer pre-treatment periods for constructing the synthetic counterfactual.

7 Discussion and Conclusion

We provide causal evidence that the *process* of negotiation – beyond the price it generates – affects whether transactions occur. Imposing fixed pricing for one model at a multi-model dealership reduced sales by 8.5 percent overall and 11.8 percent for the targeted model, even when the fixed price was below the prior negotiated average. Within-dealer demand switching – concentrated among salespeople most reliant on price discussion – points to an agent-side incentive channel rather than the price level or uniform consumer discomfort. The existing pricing-regime literature evaluates negotiation through its price outcomes; our evidence suggests that the sales process, and the people who execute it, belong in the framework.

In markets organized around personal selling, negotiation is not just a price-generation mechanism – it is a selling tool, and removing it without expanding the salesperson's alternative toolkit risks net sales losses that offset the intended gains from price transparency. Because the tool removal is independent from whether the fixed price was above or below the prior average, firms cannot solve the problem by setting the "right" price. The negotiation process provides value beyond the specific price it generates.

We acknowledge several limitations of our findings. First, we study a single model at a single manufacturer; generalization to luxury segments, weaker personal-selling cultures, or markets where digital tools have displaced much of the salesperson's role remains open. Second, we observe behavioral outcomes but cannot fully separate two channels through which removing negotiation may reduce sales: the loss of price discrimination and the degradation of the sales process. The salesperson heterogeneity evidence shifts weight toward the salesforce channel, but we cannot rule out that price discrimination accounts for a substantial share. Third, we capture short-run effects; whether the effect attenuates as salespeople develop alternative repertoires or compounds as the policy extends to additional models is an important open question. Fourth, one may argue that our setting is partial equilibrium: the one-price policy applies to one model while negotiation remains available for the rest of the portfolio. One may argue that demand switching would vanish once fixed pricing extends to all models, because salespeople would have nowhere to redirect customers. Two observations qualify this concern. On magnitude, the partial-equilibrium setting means our sales-decline estimate is an conservative estimate of the global effect of fully removing negotiation for the entire brand. Absent the redirection fallback, the total sales loss would likely be larger. On identification, the within-dealer, across-model pricing policy variation is precisely what allows us to isolate the agent-side channel – the salesperson's dependence on negotiation as a closing tool – from price-outcome-level and demand-side explanations. This channel applies whenever a personal-selling firm constrains the salesperson's toolkit, whether for one product or an entire portfolio.

References

Abadie A (2021) Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature* 59(2):391–425, URL `http://dx.doi.org/10.1257/jel.20191450`.

Abadie A, Diamond A, Hainmueller J (2010) Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association* 105(490):493–505, URL `http://dx.doi.org/10.1198/jasa.2009.ap08746`.

Anderson E, Oliver RL (1987) Perspectives on behavior-based versus outcome-based salesforce control systems. *Journal of Marketing* 51(4):76–88.

Arkhangelsky D, Athey S, Hirshberg DA, Imbens GW, Wager S (2021) Synthetic difference-in-differences. *American Economic Review* 111(12):4088–4118.

Ayres I, Siegelman P (1995) Race and gender discrimination in bargaining for a new car. *The American Economic Review* 85(3):304–321.

Barreca AI, Guldi M, Lindo JM, Waddell GR (2011) Saving babies? Revisiting the effect of very low birth weight classification. *The Quarterly Journal of Economics* 126(4):2117–2123.

Bester H (1993) Bargaining versus price competition in markets with quality uncertainty. *American Economic Review* 83(1):278–288.

Busse MR, Simester DI, Zettelmeyer F (2010) "the best price you'll ever get": The 2005 employee discount pricing promotions in the U.S. automobile industry. *Marketing Science* 29(2):268–290.

Chen Y, Yang S, Zhao Y (2008) A simultaneous model of consumer brand choice and negotiated price. *Management Science* 54(3):538–549.

Chetty R, Friedman JN, Stepner M, Opportunity Insights Team (2024) The economic impacts of COVID-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics* 139(2):829–889.

Chiou L, Tucker CE (2022) How do restrictions on advertising affect consumer search? *Management Science* 68(2):866–882.

Chung DJ, Steenburgh T, Sudhir K (2014) Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Science* 33(2):165–187.

Desai PS, Purohit D (2004) "let me talk to my manager": Haggling in a competitive environment. *Marketing Science* 23(2):219–233.

Gill D, Thanassoulis J (2016) Competition in posted prices with stochastic discounts. *The Economic Journal* 126(594):1528–1570.

Goldberg PK (1996) Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy* 104(3):622–654.

Goldberg SG, Johnson GA, Shriver SK (2024) Regulating privacy online: An economic evaluation of the GDPR. *American Economic Journal: Economic Policy* 16(1):325–358.

Holmström B, Milgrom P (1991) Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization* 7(Special Issue):24–52.

Jindal P, Newberry P (2018) To bargain or not to bargain: The role of fixed costs in price negotiations. *Journal of Marketing Research* 55(6):832–851.

Joseph K (2001) On the optimality of delegating pricing authority to the sales force. *Journal of Marketing* 65(1):62–70.

Lal R (1986) Technical note—delegating pricing responsibility to the salesforce. *Marketing Science* 5(2):159–168.

Larsen BJ (2021) The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *The Review of Economic Studies* 88(2):851–882.

Li J, Lim N, Chen H (2020) Examining salesperson effort allocation in teams: A randomized field experiment. *Marketing Science* 39(6):1122–1141.

Lim N, Chen H (2014) When do group incentives for salespeople work? *Journal of Marketing Research* 51(3):320–334.

Myerson RB, Satterthwaite MA (1983) Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29(2):265–281.

Riley JG, Zeckhauser RJ (1983) Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics* 98(2):267–289.

Roth J, Sant'Anna PHC, Bilinski A, Poe J (2023) What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics* 235(2):2218–2244.

Scott Morton F, Silva-Risso J, Zettelmeyer F (2011) What matters in a price negotiation: Evidence from the U.S. auto retailing industry. *Quantitative Marketing and Economics* 9(4):365–402.

Scott Morton F, Zettelmeyer F, Silva-Risso J (2001) Internet car retailing. *Journal of Industrial Economics* 49(4):501–519.

Scott Morton F, Zettelmeyer F, Silva-Risso J (2003) Consumer information and discrimination: Does the Internet affect the pricing of new cars to women and minorities? *Quantitative Marketing and Economics* 1(1):65–92.

Simester D, Zhang J (2014) Why do salespeople spend so much time lobbying for low prices? *Marketing Science* 33(6):796–808.

Wang R (1995) Bargaining versus posted-price selling. *European Economic Review* 39(9):1747–1764.

Weitz BA (1981) Effectiveness in sales interactions: A contingency framework. *Journal of Marketing* 45(1):85–103.

Weitz BA, Sujan H, Sujan M (1986) Knowledge, motivation, and adaptive behavior: A framework for improving selling effectiveness. *Journal of Marketing* 50(4):174–191.

Yang B, Chan TY, Thomadsen R (2019) A salesforce-driven model of consumer choice. *Marketing Science* 38(5):871–887.

Zeng X, Dasgupta S, Weinberg CB (2016) The competitive implications of a "no-haggle" pricing strategy when others negotiate: Findings from a natural experiment. *International Journal of Research in Marketing* 33(4):907–923.

Zettelmeyer F, Scott Morton F, Silva-Risso J (2006) How the Internet lowers prices: Evidence from matched survey and automobile transaction data. *Journal of Marketing Research* 43(2):168–181.

Tables and Figures

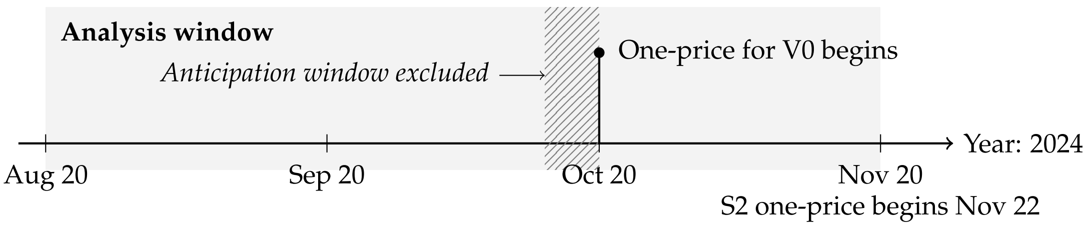

Figure 1: One-Price Policy Rollout Timeline

Table 1: Retained Car Series and Average Invoice Prices

Car Series	Avg. Invoice Price (yuan)
S1	136,776
S2	135,706
S3	128,227
V0	101,099
S4	98,533
S5	87,863
S6	85,596
S7	70,252

Notes: Average invoice prices are computed over the analysis window in both 2023 and 2024. The highlighted row (V0) is the model subject to the one-price policy.

Table 2: Summary Statistics

	Mean	SD	Min	Max	N
Panel A: Transaction Outcomes					
Total sales	6.855	7.026	<1	134	17,333
V0 sales	1.291	1.865	<1	83	17,333
Other model sales	5.564	5.991	<1	121	17,333
V0 share	0.198	0.206	0	1	17,333
Other share	0.802	0.206	0	1	17,333
Panel B: Lead and Funnel Metrics					
Lead count	6.885	8.982	<1	357	245,036
Conversion rate	0.134	0.330	0	1	48,260
Store visit rate	0.166	0.324	0	1	48,260
Decision duration (days)	6.542	10.562	<1	184	21,045

Notes: This table reports summary statistics for the estimation panel. Sales volumes, lead counts, and decision durations are rescaled by constant factors for confidentiality. Shares and rates are reported in natural units. For rescaled variables, minimum values are reported as <1 to reduce disclosure risk. All statistical analyses are conducted using the original, unscaled data. Panel A reports transaction outcomes ($N = 17{,}333$). Panel B reports lead and funnel metrics.

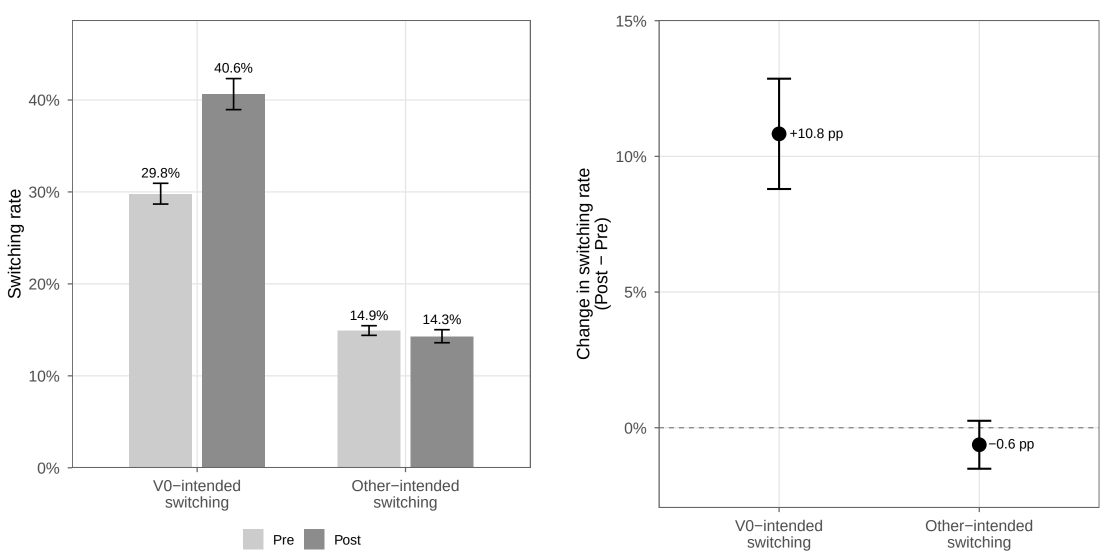

(a) Switching rate by period

(b) Change in switching rate (Post − Pre)

Figure 2: Customer Switching Patterns: Intended vs. Actual Model

Notes: Panel (a) plots the share of customers whose actual purchased model differs from their intended model, separately for the pre-policy and post-policy periods. Panel (b) plots the pre-to-post change in switching rate. 95% confidence intervals are shown in both panels.

Table 3: SDID Estimates: Effect of One-Price Policy on Sales

	Dealer level	Salesperson level
Panel A: Sales		
V0 sales (log)	-0.1254^{***}	-0.0304^{***}
	(0.0213)	(0.0038)
Total sales (log)	-0.0892^{***}	-0.0021
	(0.0339)	(0.0083)
V0 sales share (pp)	-3.255^{***}	-1.064^{***}
	(0.711)	(0.274)
Other model sales share (pp)	2.229^{**}	3.302^{***}
	(1.068)	(0.533)
Panel B: Intention Change		
Intended V0 $\rightarrow$ buy others (log)	0.0484^{***}	0.0127^{***}
	(0.0113)	(0.0017)
N (dealers/salespeople)	1,483	13,383

Notes: Each cell reports a separate SDID estimate with jackknife SE in parentheses. Dealer-level: 688 treated (2024), 795 donors (2023). Salesperson-level: 5,916 treated, 7,467 donors. Share estimates in percentage points. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

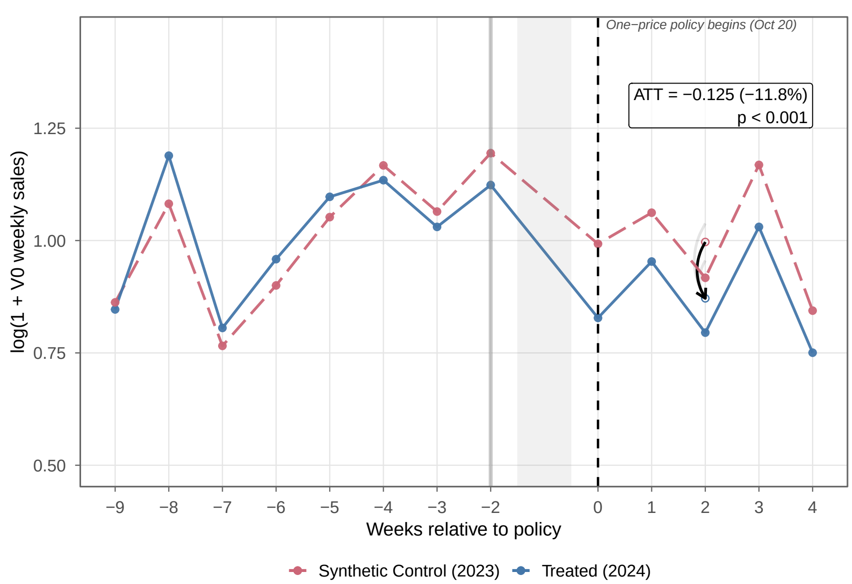

Figure 3: Pre-Treatment Fit and Post-Policy Divergence

Notes: This figure plots the treated dealers' outcome trajectory against the SDID synthetic control. The dashed vertical line marks the policy date (October 20, 2024); the vertical gray line marks the last pre-policy period included in estimation; the shaded band indicates the one-week anticipation window excluded from estimation. The post-treatment divergence is consistent with a discrete structural break at the policy date.

Table 4: Customer Heterogeneity: Age and Experience

Panel A: Customer Age

	≤ 33	34–49	≥ 50
Intended V0 $\rightarrow$ buy others (log)	0.0072	0.0290***	0.0405***
	(0.0052)	(0.0094)	(0.0069)
Other model sales share (pp)	−0.381	3.947***	6.953***
	(1.437)	(1.200)	(1.256)
N (dealers)	1,435	1,467	1,427

Panel B: Customer Experience

	Experienced	Inexperienced
Intended V0 $\rightarrow$ buy others (log)	0.0050	0.0154*
	(0.0066)	(0.0082)
Other model sales share (pp)	2.553	4.865**
	(1.777)	(2.092)
N (dealers)	708	926

Notes: Each cell reports a separate SDID estimate with jackknife standard errors in parentheses. N reports total dealer counts (treated + donor); each treated subgroup is matched to its corresponding 2023 donor pool. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5: Demand Switching: Dealer-Level SDID

	Estimate
Intended V0 $\rightarrow$ buy others (log)	0.0484***
	(0.0113)
Walk-in $\rightarrow$ buy others (log)	0.0306*
	(0.0181)
Switching rate: V0 intention (pp)	5.421***
	(1.300)
Switching rate: Other model intention (pp)	2.494
	(1.623)
Switching rate: No stated intention (pp)	2.074*
	(1.165)
N (dealers)	1,483

Notes: Each cell reports a separate SDID estimate with jackknife standard errors in parentheses. Outcomes are log-transformed counts or percentage-point shares. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6: Individual Switching: Logistic Regression with Dealer Fixed Effects

	Purchase other models	Odds ratio
Post	0.017	1.017
	(0.034)	
V0-intended	−3.279***	0.038
	(0.033)	
Walk-in	−0.931***	0.394
	(0.024)	
Post × V0-intended	0.594***	1.811
	(0.055)	
Post × Walk-in	0.152***	1.164
	(0.041)	
Dealer FE	Yes	
N (transactions)	127,693	

Notes: Logistic regression of Pr(purchase $\neq$ V0) on post-policy indicator, intention group, and their interactions, with dealer fixed effects (Equation 6.1). Standard errors in parentheses. AME standard errors in parentheses computed via dealer-clustered bootstrap (500 iterations). Sample restricted to completed purchases in the analysis window. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 7: Salesperson Heterogeneity: Reliance on Price Negotiation

Panel A: Price-First Discussion Rate		
	High reliance	Low reliance
Intended V0 → buy others (log)	0.0639***	0.0307*
	(0.0181)	(0.0178)
N (dealers)	618	606
Panel B: Sales/Customer Price Mention Ratio		
	Higher ratio	Lower ratio
Intended V0 → buy others (log)	0.0508**	0.0383**
	(0.0200)	(0.0171)
N (dealers)	614	609

Notes: Each cell reports a separate SDID estimate with jackknife SEs in parentheses. N reports total units (treated + donor); each treated subgroup is matched to its corresponding 2023 donor pool. Panel A: dealers classified by pre-policy average salesperson price-first discussion rate (median split at 48.6%). Panel B: dealers classified by pre-policy average sales/customer price mention ratio (median split at 1.461). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 8: Other-Model Mentions by Salesperson Reliance on Price Negotiation

Panel A: Post-policy change in other-model mentions by SA reliance			
	Pre (%)	Post (%)	Diff (pp)
High reliance	25.251	28.389	3.139**
	(0.888)	(1.085)	(1.402)
Low reliance	25.821	24.331	−1.490
	(0.860)	(0.964)	(1.292)

Panel B: Reliance interaction	
	Other-model mention (pp)
Post	−1.397
	(1.383)
Post × High reliance	5.427**
	(2.164)
Dealer FE	Yes
N (leads)	8,649

Notes: Panel A: within-group proportions tests. Panel B: linear probability model with dealer fixed effects and dealer-clustered standard errors; coefficients in percentage points. Sample: V0-intended leads with transcripts (>50 characters) in the analysis window. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 9: Price-Level Test: Above vs. Below Fixed Price

Policy impact:	Higher Pre-Policy Price *(V0 became cheaper)*	Lower Pre-Policy Price *(V0 became more expensive)*
Panel A: Sales		
V0 sales (log)	−0.1404***	−0.1607***
	(0.0280)	(0.0526)
Total sales (log)	−0.0877**	−0.1070
	(0.0410)	(0.0668)
V0 sales share (pp)	−4.561***	−3.648**
	(0.951)	(1.741)
Other model sales share (pp)	−0.036	0.996
	(1.458)	(2.174)
Panel B: Intention Change		
Intended V0 → buy others (log)	0.0380***	0.0592*
	(0.0142)	(0.0330)
N (dealers)	1,042	278

Notes: Each cell reports a separate SDID estimate with jackknife standard errors in parentheses. Dealers classified by pre-policy average V0 final user price relative to the one-price. "Higher" dealers experienced a price decrease under the policy; "Lower" dealers experienced a price increase. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 10: Lead Quality Around the One-Price Policy

	V0	Other
Conversion rate	0.004	0.001
	(0.005)	(0.004)
Decision duration (days)	0.942	−0.744
	(0.845)	(1.022)
Store visit rate	0.007	−0.002
	(0.006)	(0.006)
N (dealers)	1,340	1,405

Notes: SDID with jackknife standard errors in parentheses. N reports total dealer counts (treated + donor); each treated unit is matched to corresponding 2023 donor dealers. All outcomes are in levels. "V0" filters to V0-intended leads; "Other" filters to non-V0-intended leads. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Online Appendix

A Variable Construction

Customer Experience. We classify customers as “experienced” (= 1) or “inexperienced” (= 0) based on LLM analysis of their conversation transcripts. The classification uses a two-stage evaluation procedure, with the prompt below. In Stage 1, the model identifies indicators of purchase experience from the transcript text. In Stage 2, the model synthesizes these indicators into a binary classification. The classification is performed by Gemini 2.5 Flash Lite (temperature = 0.1, max output = 256 tokens). The full system prompt (translated from Chinese) is reproduced below.

System Prompt – Customer Experience Classification

```
You are an experienced car sales analyst.  Read TWO conversation stages for one
shopper and decide whether the shopper is EXPERIENCED (binary).  Put extra weight
on the second-stage conversation.  Use the numeric context provided as weak evidence
only.  Return JSON with the required keys - no extra text.
```

Strong positive indicators (any 2 suffice):

1. *Model familiarity* – mentions specific models, trims, model years, or configurations and makes meaningful comparisons across brands.
2. *Price concept fluency* – correctly uses or calculates guide price, bare-car price, transaction price, landing price, discount percentages; can state a target landing price or compare quotes across dealers.
3. *Policy/promotion knowledge* – references one-price policy, manufacturer subsidies, purchase-tax promotions, trade-in subsidies, warranty/extended warranty, deposit/cancellation terms, or production/delivery timelines.
4. *Financing literacy* – understands down payment, monthly payment, APR, 36/48/60-month terms, residual value, financial lease, early payoff penalties; can compare bank vs. dealer financing or request removal of financial service fees.
5. *Bargaining structure* – requests itemized quotes, asks to remove add-ons, demands the dealer match a competitor's offer, or cites third-party platforms or real-world quotes.
6. *Technical detail* – unprompted correct use of technical terminology (e.g., 2.0L NA vs. 1.5T, China VI-B, L2 ADAS, OTA).

Strong negative indicators (lean toward 0, especially in Stage 2): basic-concept questions; explicit first-time-buyer statements; listing $\geq$ 3 models with no comparison logic; dominant affordability distress with no substantive information.

Numeric context (weak evidence; ignore if null): weakly positive if decision days $\leq$ average, visit count $\leq$ average, or age $\geq$ average; weakly negative otherwise.

Stage weighting: Stage 2 weight = 0.7; Stage 1 weight = 0.3. Conflicts resolved in favor of Stage 2.

Decision rules: (1) Stage 2 has $\geq$ 2 strong positives $\Rightarrow$ experienced = 1. (2) Stage 2 has exactly 1 strong positive and Stage 1 has $\geq$ 1, or numeric weak-positives $\geq$ 2 $\Rightarrow$ experienced = 1. (3) Stage 2 has $\geq$ 2 strong negatives with no clear positive $\Rightarrow$ experienced = 0. (4) Otherwise: judge holistically from Stage 2; if evidence is thin or contradictory, default to 0.

Output: `{"experienced": 0 or 1}`

Price-Discussion Frequency. For each salesperson, we compute the share of pre-policy conversations in which the salesperson initiated the price topic before the customer. This variable – the *price-first rate* – is derived from textual analysis of call transcripts that identifies the first party to raise pricing, discounts, or transaction terms. The dealer-level classification variable used in Table 7 is the average price-first rate across all salespeople at each dealer, split at the sample median.

Price Mention Ratio. We compute the ratio of salesperson price mentions to customer price mentions in each conversation. Price mentions are identified via keyword matching in the call transcript text for terms related to pricing, discounts, and transaction amounts. The dealer-level classification variable used in Table 7, Panel B is the average mention ratio across all pre-policy conversations at each dealer, split at the sample median.

V0-to-Other Switching. A binary indicator equal to one if the customer's intended car series (as recorded in the lead system at initial inquiry) is a V0 variant but the actual purchased car series is not.

No-Haggle Language and Other-Model Redirect. We construct two binary indicators at the lead level from call transcripts (Table 8). Transcripts are formed by concatenating first-stage and second-stage recordings; only salesperson utterances are retained.
Haggle-refusal language equals one if the salesperson's text contains any of approximately 11 phrases in which the salesperson refuses to negotiate on the focal model's price. Keywords include: cannot discount (不能优惠), no bargaining (不讲价, 不还价), unified price (统一价, 全国统一), fixed price (固定价), and lowest already (最低了, 不能再低).
Other-model mention equals one if the salesperson names another model in the brand's lineup by its model name or uses indirect redirect phrases (e.g., look at another one, also have a model, more cost-effective). We match six specific model names and approximately 15 redirect phrases. The sample consists of 8,720 V0-intended leads with transcripts exceeding 50 characters, covering the analysis window.

B Supplementary Results

Table B.1: SDID Unit Weights and Time Weights

Panel A: Unit Weights over 2023 Donor Dealers ($\hat{\omega}_i$)		
	Total Sales	V0 Sales
Mean	0.0013	0.0013
Median	0.0012	0.0013
Max	0.0035	0.0025
$N_{\omega>0.01}$	0	0
Effective N	689	746
Panel B: Time Weights ($\hat{\lambda}_t$)		
Pre-Treatment Period	Total Sales	V0 Sales
Week -8	0.020	0.113
Week -7	0.000	0.059
Week -6	0.112	0.126
Week -5	0.150	0.134
Week -4	0.104	0.120
Week -3	0.128	0.151
Week -2	0.125	0.088
Week -1	0.361	0.209
Pre-treatment RMSE	0.102	0.056

Notes: Panel A reports the distribution of SDID unit weights ($\hat{\omega}_i$) over the 795 donor dealers (2023 observations). Effective $N = 1/\sum \omega_i^2$ measures the effective number of donors. Panel B reports SDID time weights ($\hat{\lambda}_t$) for each pre-treatment period. Pre-treatment RMSE measures the fit of the synthetic control in the pre-treatment period. $N_{\text{donors}} = 795$, $N_{\text{treated}} = 688$.

Table B.2: In-Space Permutation Test

	Actual	Perm. mean	Perm. SD	$\lvert$Placebo$\rvert \geq \lvert$Actual$\rvert$	p-value
V0 sales (log)	-0.1254	0.0002	0.0224	0	0.002
Total sales (log)	-0.0892	-0.0000	0.0270	0	0.002
V0 share (pp)	-3.2549	-0.0346	0.5764	0	0.002
Other share (pp)	2.2293	0.0390	0.9529	10	0.022
V0 to other (log)	0.0484	0.0003	0.0096	0	0.002

Notes: In-space permutation test following Abadie et al. (2010). Treatment and donor labels are randomly reassigned among all units (keeping the same number of treated and donors), and SDID is re-estimated. Repeated 500 times per outcome. Two-sided p-value: $(\#|\hat{\tau}_{\text{placebo}}| \geq |\hat{\tau}_{\text{actual}}| + 1)/(N_{\text{perms}} + 1)$.

Table B.3: Within-Year SDID: V0 vs. Premium Controls

	Estimate
Panel A: Sales	
V0 sales (log)	-0.0979^{***}
	(0.0187)
Total sales (log)	0.0156
	(0.0240)
V0 sales share (pp)	-1.475^{***}
	(0.524)
Other model sales share (pp)	1.876^{**}
	(0.816)
Panel B: Intention Change	
Intended V0 $\rightarrow$ buy others (log)	0.0812^{***}
	(0.0178)
N (dealer $\times$ model)	21,190

Notes: Each cell reports a separate SDID estimate with jackknife SE in parentheses. Within-year design using premium models (P1, P2, P3) as the control group. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.4: Anticipation Window Robustness

	Estimate
Panel A: Sales	
V0 sales (log)	-0.1491^{***}
	(0.0221)
Total sales (log)	-0.0674^{**}
	(0.0340)
V0 sales share (pp)	-4.060^{***}
	(0.752)
Other model sales share (pp)	2.996^{**}
	(1.174)
Panel B: Intention Change	
Intended V0 $\rightarrow$ buy others (log)	0.0587^{***}
	(0.0111)
N (dealers)	1,483

Notes: Each cell reports a separate SDID estimate with jackknife SE in parentheses. Treatment cutoff shifted to Oct 13 (anticipation start); no gap week dropped. Sample: 688 treated dealers (2024) and 795 donor dealers (2023). Share estimates in percentage points. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.5: Window Sensitivity: SDID Estimates Across Symmetric Windows

Window (weeks)	Estimate
4	-0.052^{**}
	(0.026)
5	-0.092^{***}
	(0.022)
6	-0.122^{***}
	(0.023)
Main spec	-0.1254^{***}

Notes: Each row reports a separate SDID estimate at the indicated window. This table re-estimates the main SDID specification for V0 sales using symmetric pre- and post-treatment windows around the policy date. Standard errors in parentheses, computed via jackknife. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.6: Salesperson Heterogeneity: Pre-Policy Balance

	High rel.	Low rel.	Diff (p)
Price-first rate (%)†	63.85	32.16	31.69 ($<$ 0.001)
Total conversations	15.41	15.65	-0.24 (0.845)
N (salespeople)	6.55	6.57	-0.02 (0.964)

Notes: Pre-policy group means with two-sample t-test p-values. Total conversations and N salespeople are balanced ($p > 0.8$).
†Price-first rate is the classification variable; the significant difference is by construction.

Table B.7: Salesperson Heterogeneity: Alternative Classifications

	High reliance	Low reliance
Panel A: Top vs. bottom 25% of dealers		
Intended V0 $\rightarrow$ buy others (log)	0.0759^{***}	0.0136
	(0.0262)	(0.0245)
N (dealers)	307	300
Panel B: Salesperson-level classification		
Intended V0 $\rightarrow$ buy others (log)	0.0161^{***}	0.0120^{***}
	(0.0041)	(0.0036)
N (salespeople)	2,271	3,143

Notes: Each cell reports a separate SDID estimate with jackknife SEs in parentheses. N reports total units (treated + donor); each treated subgroup is matched to its corresponding 2023 donor pool. Panel A compares dealers in the top 25% (high reliance) and bottom 25% (low reliance) of the pre-policy price-first discussion rate, excluding dealers near the median margin whose reliance classification and predicted behavioral response are less clear. Panel B is a salesperson-level analysis that classifies each salesperson by their own pre-policy price-first rate with a median split. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

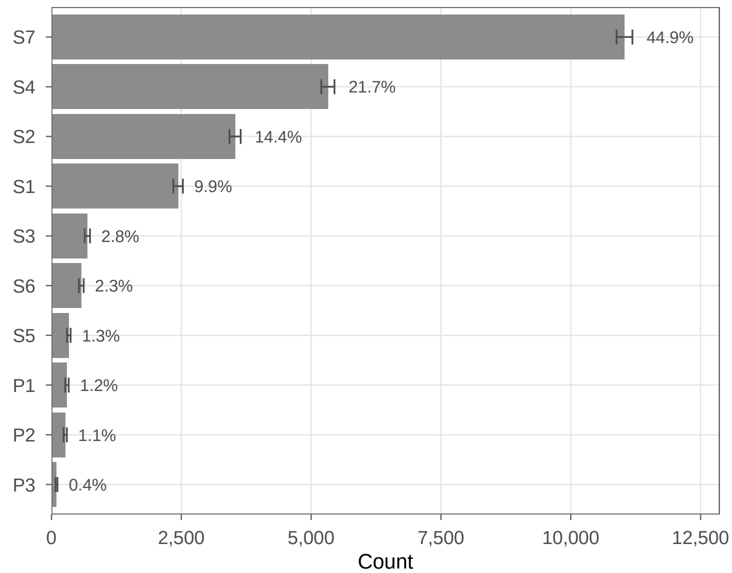

Figure B.1: Switching Destinations from V0

Notes: This figure shows the destination models for customers who originally intended to purchase V0 but ultimately purchased a different model. Bars report the count of switchers to each target car series; 95% Wilson confidence intervals are shown.

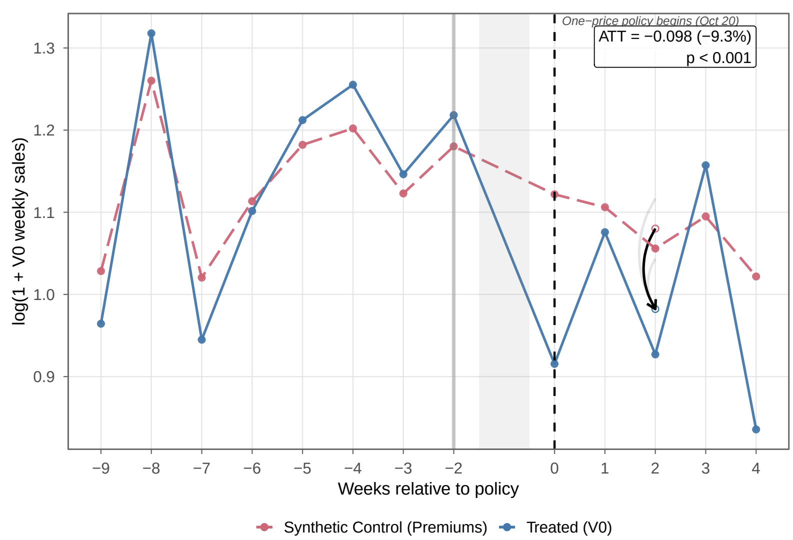

Figure B.2: Within-Year SDID: V0 vs. Premium Controls

Notes: This figure plots V0's outcome trajectory against the SDID synthetic control constructed from premium models (P1, P2, P3) within the 2024 analysis window. The dashed vertical line marks the policy date; the vertical gray line marks the last pre-policy period included in estimation; the shaded band indicates the one-week anticipation window excluded from estimation.